

Abstract

A database system for selecting rows from a partitioned database table is disclosed. The partitioned database table includes rows and columns and is divided into partitions with at least one of the partitions in the table being populated by one or more rows. The system includes one or more nodes, each of the one or more nodes providing access to one or more CPUs. Each of the one or more CPUs provides access to one or more virtual processes. Each process is configured to manage data, including the partitioned database table, stored in one of a plurality of data-storage facilities. The system also includes a partitioned table access component configured to select rows from the table by creating a file context, which stores at least location data for a row and a first value associated with the row, for each populated partition, determining the lowest first value stored by the file contexts, and identifying rows with a particular first value by reading the file contexts.